

### (12) United States Patent

Szabo et al.

# (10) **Patent No.:**

US 9,338,095 B2

(45) Date of Patent:

May 10, 2016

#### (54) DATA FLOW SEGMENT OPTIMIZED FOR **HOT FLOWS**

(71) Applicant: **F5 NETWORKS, INC.**, Seattle, WA

(US)

Inventors: Paul Imre Szabo, Shoreline, WA (US);

Peter Michael Thornewell, Seattle, WA

(US); Timothy Scott Michels, Greenacres, WA (US)

(73) Assignee: F5 Networks, Inc., Seattle, WA (US)

(\*) Notice: Subject to any disclaimer, the term of this

patent is extended or adjusted under 35

U.S.C. 154(b) by 304 days.

Appl. No.: 13/802,254

(22)Filed: Mar. 13, 2013

(65)**Prior Publication Data** 

> US 2013/0294239 A1 Nov. 7, 2013

### Related U.S. Application Data

- (60) Provisional application No. 61/641,251, filed on May 1, 2012.
- (51) Int. Cl. H04L 12/26 (2006.01)H04L 12/801 (2013.01)H04L 12/721 (2013.01)H04L 12/715 (2013.01)H04L 29/08 (2006.01)
- (52) U.S. Cl.

CPC ...... H04L 47/10 (2013.01); H04L 45/38 (2013.01); H04L 45/64 (2013.01); H04L 47/12 (2013.01); H04L 67/1002 (2013.01)

(58) Field of Classification Search

None

See application file for complete search history.

#### (56)References Cited

#### U.S. PATENT DOCUMENTS

3,950,735 A 4/1976 Patel 2/1987 George et al. 4,644,532 A 4,965,772 A 10/1990 Daniel et al. (Continued)

#### FOREIGN PATENT DOCUMENTS

EP0 744 850 A2 11/1996 WO 91/14326 A2 9/1991

(Continued) OTHER PUBLICATIONS

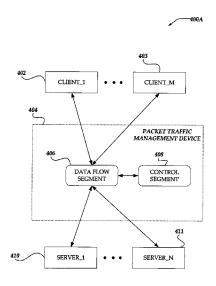
Office Communication for U.S. Appl. No. 13/802,169 mailed on Oct. 9, 2004.

(Continued)

Primary Examiner — Jutai Kao (74) Attorney, Agent, or Firm — John W. Branch; Lowe Graham Jones PLLC ABSTRACT

Embodiments are directed towards improving the performance of network traffic management devices by optimizing the management of hot connection flows. A packet traffic management device ("PTMD") may employ a data flow segment ("DFS") and control segment ("CS"). The CS may perform high-level control functions and per-flow policy enforcement for connection flows maintained at the DFS, while the DFS may perform statistics gathering, per-packet policy enforcement (e.g., packet address translations), or the like, on connection flows maintained at the DFS. The DFS may include high-speed flow caches and other high-speed components that may be comprised of high-performance computer memory. Making efficient use of the high speed flow cache capacity may be improved by maximizing the number of hot connection flows and minimizing the number of malicious and/or in-operative connections flows (e.g., nongenuine flows) that may have flow control data stored in the high-speed flow cache.

#### 21 Claims, 12 Drawing Sheets



(56)	(56) References Cited		2002/0138		9/2002	
U.S. PATENT DOCUMENTS			2004/0039 2004/0049	596 A1	3/2004	Colby et al. Schuehler et al.
5.000.006	6/1001	D . 1	2004/0111 2006/0095			Boivie et al. Van Doren et al.
5,023,826 A 5,053,953 A			2008/0162		7/2008	Kapoor et al.
5,299,312 A		Rocco, Jr.	2008/0181	226 A1	7/2008	Varier et al.
5,327,529 A		Fults et al.	2008/0256 2009/0003			Gilde et al. Okholm et al.
5,367,635 A 5,371,852 A		Bauer et al. Attanasio et al.	2009/0003			Chen et al.
5,406,502 A		Haramaty et al.	2009/0209	262 A1	8/2009	Stamoulis et al.
5,475,857 A	12/1995	Dally				Foschiano et al 709/233
5,517,617 A		Sathaye et al.	2010/0121 2010/0315			Samuels et al. Turanyi
5,519,694 A 5,519,778 A		Brewer et al. Leighton et al.	2011/0075		3/2011	Koodli et al 370/401
5,521,591 A	5/1996	Arora et al.	2011/0179			Lindsay
5,528,701 A		Aref Fitzgerald et al.	2012/0320 2013/0044			Venkataramanan et al. Lappetelainen et al.
5,581,764 A 5,596,742 A		Agarwal et al.	2013/0083			Gupta et al 370/235
5,606,665 A	2/1997	Yang et al.	2014/0036	661 A1	2/2014	Campbell
5,611,049 A				FOREIGN PATENT DOCUMENTS		
5,663,018 A 5,752,023 A		Cummings et al. Choucri et al.		FOREIC	IN PALE	NI DOCUMENTS
5,761,484 A	6/1998	Agarwal et al.	WO	95/0:	5712 A2	2/1995
5,768,423 A	6/1998	Aref et al.	WO		9805 A1	3/1997
5,774,660 A 5,790,554 A		Brendel et al. Pitcher et al.	WO		5800 A1	12/1997
5,802,052 A		Venkataraman	WO WO		5829 A1 5913 A1	2/1999 2/1999
5,875,296 A	2/1999	Shi et al.	WO	99/1	0858 A2	3/1999
5,892,914 A 5,892,932 A			WO		9373 A2	8/1999
5,919,247 A		Van Hoff et al.	WO WO		4967 A1 4422 A2	12/1999 1/2000
5,936,939 A	8/1999	Des Jardins et al.	WO		1458 A1	1/2000
5,946,690 A	8/1999	Pitts		OT	HER DIT	BLICATIONS
5,949,885 A 5,951,694 A	9/1999	Leighton Choquier et al		OI	IILKI O.	BLICATIONS
5,959,990 A		Frantz et al.	Office Com	munication	n for U.S.	Appl. No. 13/461,675 mailed on
5,974,460 A	10/1999	Maddalozzo, Jr. et al.	Aug. 14, 20			137 42/552 434 11 7 1
5,983,281 A 6,006,260 A		Ogle et al. Barrick, Jr. et al.	25, 2014.	munication	for U.S. A	Appl. No. 13/772,194 mailed on Jul.
6,006,264 A		Colby et al.		Advanced 1	Encryption	n Standard (AES), Nov. 26, 2001,
6,026,452 A			NIST, all pa			
6,028,857 A 6,051,169 A		Poor Brown et al.		nmunicatio	on for U.S	. Appl. No. 13/461,675 mailed Jan.
6,078,956 A	6/2000	Bryant et al.	27, 2014.		etic A	1 N 12/772 104 1-1 I
6,085,234 A		Pitts et al.	26, 2015 (2)		. Ior U.S. A	Appl. No. 13/772,194 mailed on Jan.
6,092,196 A 6,108,703 A		Reiche Leighton et al.			for U.S. A	appl. No. 13/772,194 mailed on Apr.
6,111,876 A	8/2000	Frantz et al.	15, 2015 (9	pages).		
6,178,423 B		Douceur et al.			for U.S. A	Appl. No. 13/802,169 mailed on Feb.
6,182,139 B 6,192,051 B	31 1/2001 31 2/2001	Brendel Lipman et al.	5, 2015 (10		za Poutino	of Servlet Content to Transcoding
6,246,684 B		Chapman et al.				422124, IBM Corporation, pp. 889-
6,253,230 B		Couland et al.	890, Jun. 19			, ,, ,, , <sub>F</sub> , <sub>F</sub>
6,263,368 B 6,278,995 B		Martin Hawkinson				ication Process With Single Sign-
6,327,622 B		Jindal et al.		ch Disclos	ure 42912	28, IBM Corporation, pp. 163-164,
6,374,300 B	2 4/2002	Masters	Jan. 2000. "Transmissi	on Control	Protocol	"Wikipedia, the free encyclopedia,
6,396,833 B 6,601,084 B		Zhang et al. Bhaskaran et al.				nsmission_Control_Protocol, pp.
6,636,894 B		Short et al.	1-18, last ac	cessed Au	g. 1, 2012	•
6,650,641 B		Albert et al.				Written Opinion for International
6,742,045 B 6,751,663 B		Jordan et al. Farrell et al.				2013/038168 mailed Aug. 14, 2013. appl. No. 13/802,169 mailed on May
6,754,228 B	6/2004	Ludwig	27, 2015.	numeanon	101 U.S. A	appr. No. 15/802,109 maned on Way
6,760,775 B	7/2004	Anerousis et al.		munication	for U.S. A	appl. No. 13/802,331 mailed on Jun.
6,772,219 B 6,779,039 B		Shobatake Bommareddy et al.	5, 2015.			
6,781,986 B	8/2004	Sabaa et al.			for U.S. A	Appl. No. 13/772,194 mailed on Jul.
6,798,777 B		Ferguson et al.	28, 2015 (7 Office Com		for IIS	Appl. No. 13/461,675 mailed on
6,868,082 B 6,876,629 B		Allen, Jr. et al. Beshai et al.	Nov. 5, 201:	5 (37 page:	s).	11ppi. 110. 15/101,075 maned on
6,876,654 B		Hegde	F5 Network	s, Inc., "I	MOS Ma	nagement Guide for BIG-IP Sys-
6,888,836 B	5/2005	Cherkasova				o/en-us/products/big-ip_ltm/manu-
7,343,413 B		Gilde et al.				_guide101.html, publication
7,561,517 B 8,024,483 B		Klinker et al. Rothstein et al.	date Jan. 20	, ∠011, acc	essea on I	Dec. 14, 2015 (520 pages).
2001/0037387 A		Gilde et al.	* cited by	examiner		
			,			

<sup>\*</sup> cited by examiner

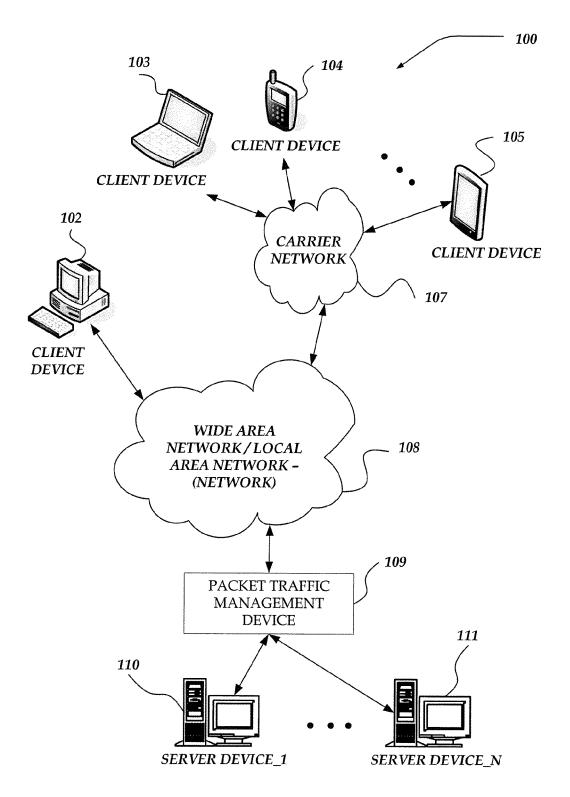


FIG. 1

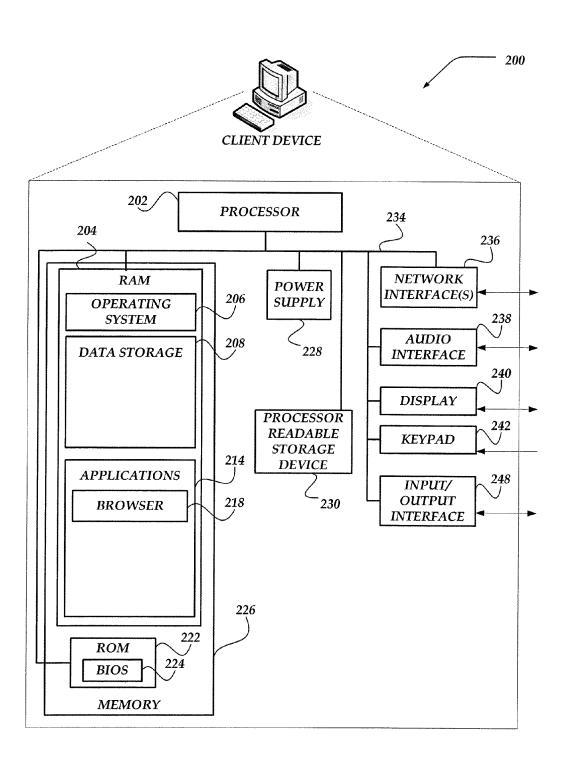
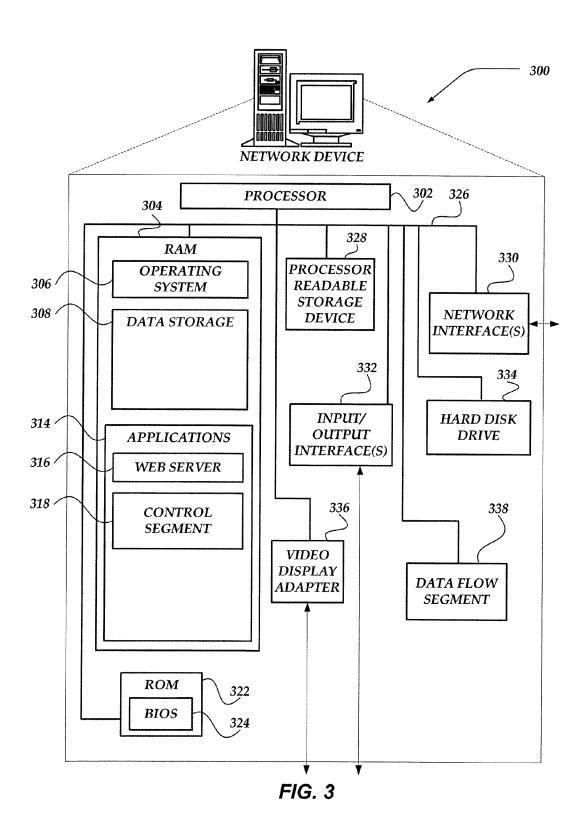


FIG. 2



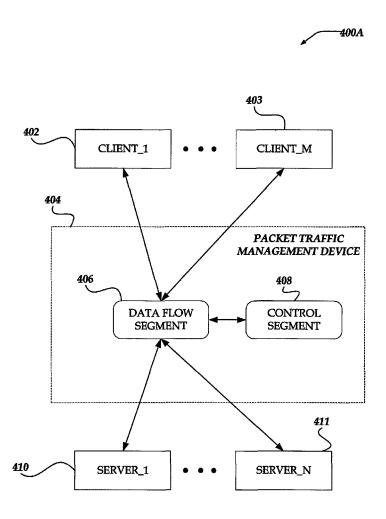


FIG. 4A

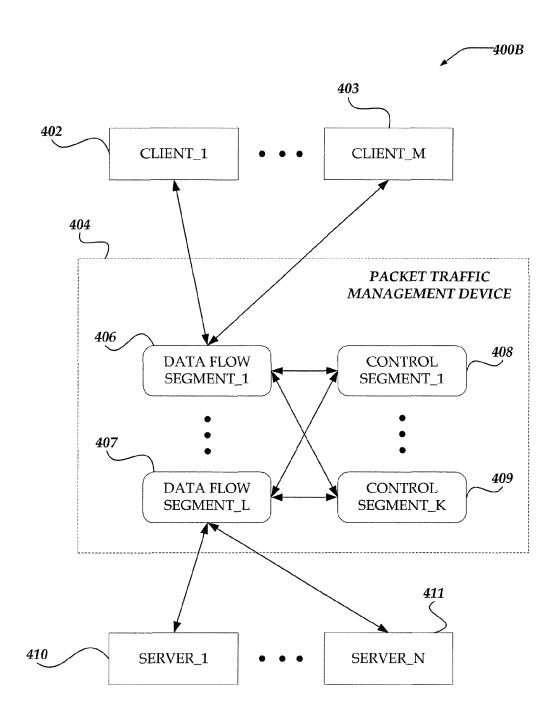


FIG. 4B

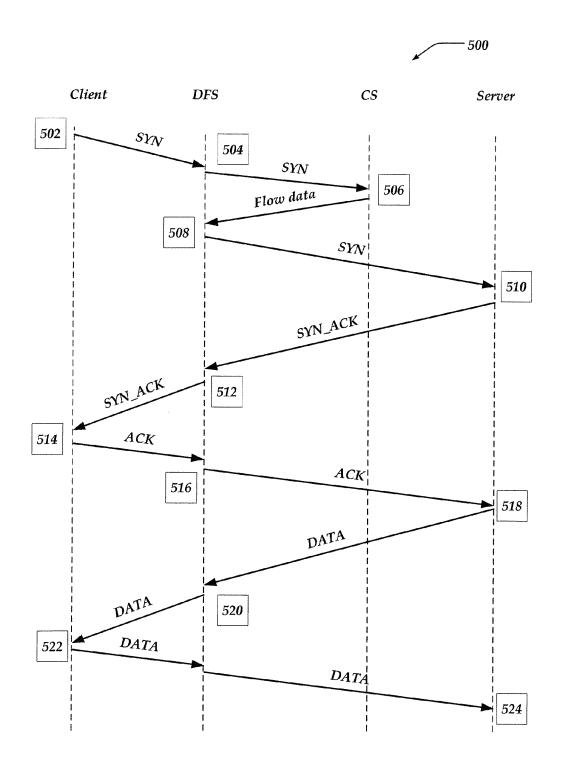


FIG. 5

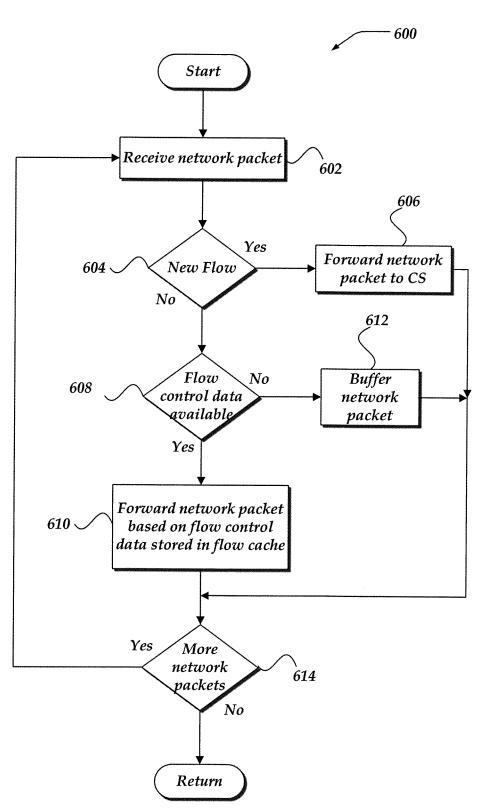


FIG. 6

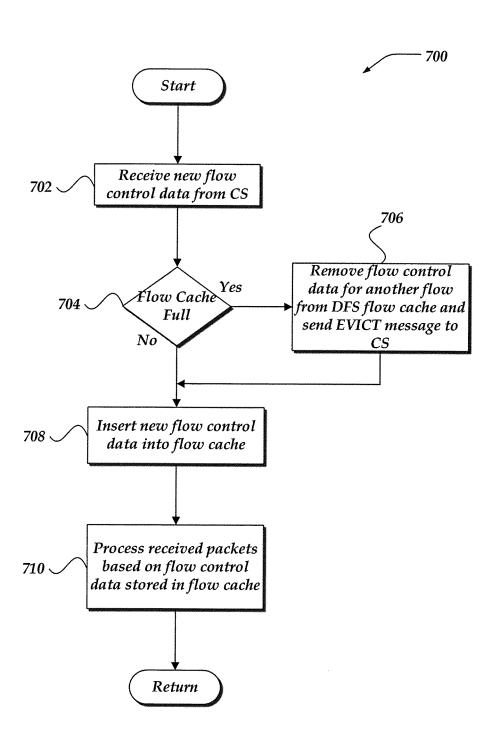


FIG. 7

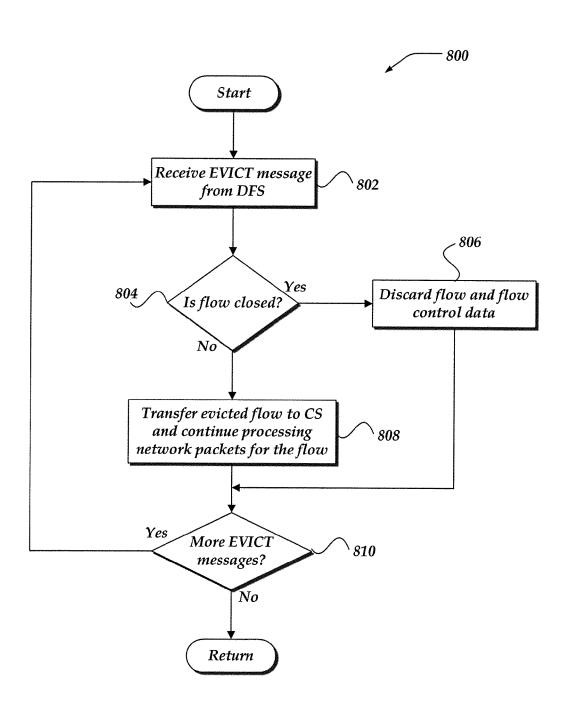


FIG. 8

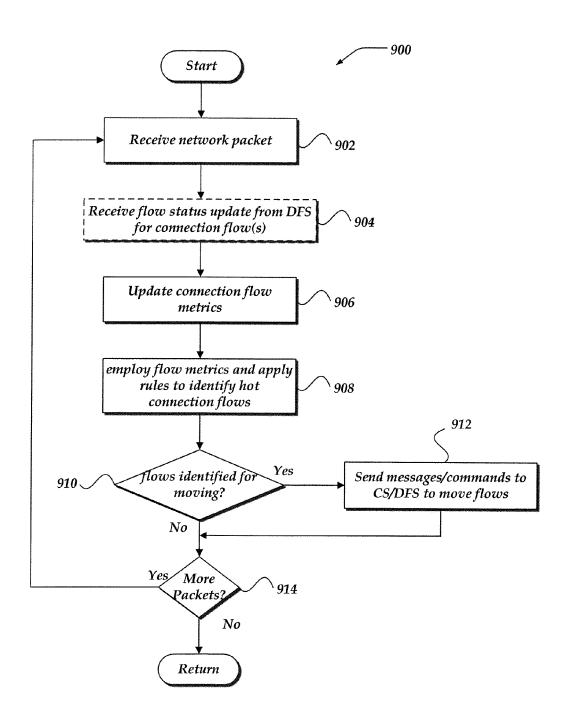


FIG. 9

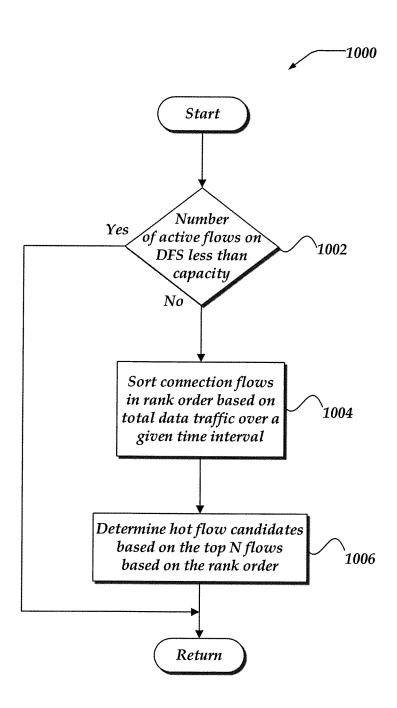


FIG. 10

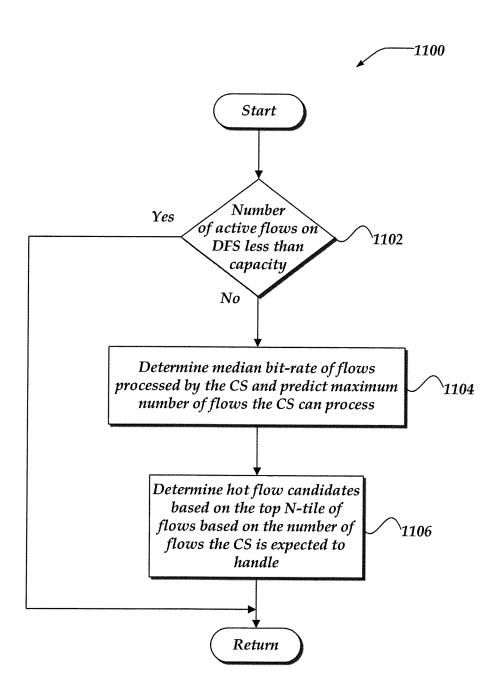


FIG. 11

1

## DATA FLOW SEGMENT OPTIMIZED FOR HOT FLOWS

#### RELATED APPLICATIONS

This application is a Utility Patent application based on a previously filed U.S. Provisional Patent application, U.S. Ser. No. 61/641,251 filed on May 1, 2012, the benefit of the filing date of which is hereby claimed under 35 U.S.C. §119(e).

#### TECHNICAL FIELD

The present invention relates generally to packet traffic management and, more particularly, but not exclusively to determining if network connection flow control data should be off-loaded to data flow segment stored in a high-speed cache.

#### BACKGROUND

The expanded use of the Internet has increased communication connections between client devices and server devices. Often, a client device establishes a network connection with a server device by using well-known protocols, such as Transmission Control Protocol/Internet Protocol ("TCP/IP"), User 25 Datagram Protocol ("UDP"), and the like. This network connection may be identified by one characteristic or a combination of characteristics, such as a source port, a destination port, a source address, a destination address, a protocol, and the like. Typically, the source address, destination address, 30 destination port, and protocol are relatively fixed for a network connection between a client device and a server device. Thus, the source port may be utilized to uniquely identify a connection between the client device and the server device. Additionally, the expansion of the Internet has led to improvements in packet traffic management. One such advancement is to split operations between a control segment and a data flow segment as described in more detail in U.S. Pat. No. 7,343,413, filed Mar. 21, 2001, and entitled "Method and System for Optimizing a Network by Independently Scaling 40 Control Segments and Data Flow," which is hereby incorporated by reference in its entirety into this patent application. Thus, it is with respect to these considerations and others that the invention has been made.

#### BRIEF DESCRIPTION OF THE DRAWINGS

Non-limiting and non-exhaustive embodiments of the present invention are described with reference to the following drawings. In the drawings, like reference numerals refer to like parts throughout the various figures unless otherwise specified. For a better understanding of the present invention, reference will be made to the following Detailed Description, which is to be read in association with the accompanying drawings, wherein:

- FIG. 1 is a system diagram of an environment in which embodiments of the invention may be implemented;
- FIG. 2 shows an embodiment of a client device that may be included in a system such as that shown in FIG. 1;
- FIG. 3 shows an embodiment of a network device that may 60 be included in a system such as that shown in FIG. 1;
- FIG. 4A and 4B illustrate overview system diagrams generally showing embodiments of a packet traffic management device disposed between client devices and server devices in accordance with the embodiments;
- FIG. 5 illustrates a sequence diagram generally showing one embodiment of a sequence for terminating a connection

2

flow at a data flow segment and establishing a new connection flow at the data flow segment in accordance with the embodiments:

- FIG. 6 shows a flowchart showing a process for packet traffic management in accordance with at least one of the various embodiments;
- FIG. 7 shows a flowchart of a process for handling new connection flows at a data flow segment in accordance with at least one of the various embodiments;
- FIG. 8 shows a flowchart of a process for handling eviction messages at a control segment in accordance with at least one of the various embodiments;
- FIG. 9 shows a flowchart of a process for determining if connection flows may be candidates for off-loading to the data flow segment in accordance with at least one of the various embodiments; and

FIGS. 10 and 11 show flowcharts of processes for identifying hot connection flows in accordance with at least one of the various embodiments.

#### DETAILED DESCRIPTION

Throughout the specification and claims, the following terms take the meanings explicitly associated herein, unless the context clearly dictates otherwise. The phrase "in one embodiment" as used herein does not necessarily refer to the same embodiment, though it may. Furthermore, the phrase "in another embodiment" as used herein does not necessarily refer to a different embodiment, although it may. Thus, as described below, various embodiments of the invention may be readily combined, without departing from the scope or spirit of the invention.

In addition, as used herein, the term "or" is an inclusive "or" operator, and is equivalent to the term "and/or," unless the context clearly dictates otherwise. The term "based on" is not exclusive and allows for being based on additional factors not described, unless the context clearly dictates otherwise. In addition, throughout the specification, the meaning of "a," "an," and "the" include plural references. The meaning of "in" includes "in" and "on."

As used herein, the term "SYN" refers to a packet transmitted utilizing TCP that includes a set synchronize control flag in a TCP header of the packet.

As used herein, the term "ACK" refers to a packet trans-45 mitted utilizing TCP that includes a set acknowledgment flag in a TCP header of the packet.

As used herein, the term "SYN\_ACK" refers to a packet transmitted utilizing TCP that includes a set synchronize control flag and a set acknowledgment flag in a TCP header of the packet.

As used herein, the term "FIN" refers to a packet transmitted utilizing TCP that includes a set no more data from sender flag in a TCP header of the packet.

As used herein, the term "FIN\_ACK" refers to a packet transmitted utilizing TCP that includes a set no more data from sender flag and a set acknowledgment flag in a TCP header of the packet. FIN\_ACK compress a FIN and ACK into one TCP packet.

As used herein, the term "tuple" refers to a set of values that identify a source and destination of a connection. In one embodiment, a 5 tuple may include a source address, a destination address, a source port, a destination port, and a protocol identifier. In at least one of the various embodiments, tuples may be used to identify network flows (e.g., connection flows).

As used herein, the terms "network flow," "connection flow,", "flow" refer to a network session that may be estab027,220,0322

lished between two endpoints. In at least one of the various embodiments, a tuple may describe the flow. In at least one of the various embodiments, flow control data associated with connection flows may be used to ensure that the network packets sent between the endpoints of a connection flow may be routed along the same path. In at least one of the various embodiments, the performance of connection oriented network protocols such as TCP/IP may impaired if network packets may be routed using varying paths and/or directed different endpoints.

3

As used herein, the term "genuine connection flow," refers to a connection flow that may have been determined to be associated with an operative client-server communication session. In contrast, a non-genuine connection flow may be associated with a malicious attack such as a SYN flood attack. In at least one of the various embodiments, characteristics a genuine connection flows may include, TCP/IP handshaking complete, evidence of bi-directional network packet exchange, or the like. Likewise, evidence that a connection flow may be non-genuine may include, half-open connections (incomplete handshaking and connection setup), few if any network packets exchanged, or the like.

As used herein, the term "hot connection flow," refers to a connection flow that may have been determined to be a candidate for off loading to a data flow segment. Hot flow connections may have characteristics such high-bandwidth utilization, quality of service priority, or the like.

As used herein, the term "high speed flow cache" refers to memory based cache used for storing flow control data that corresponds to connection flows. The cache may be accessible using, dedicated busses that may provide very fast performance based on a combination of factors that may include, wide-busses, fast clock speeds, dedicated channels, specialized read and/or write buffer, hardware proximity, temperature control, or the like. Also, the high speed flow cache may 35 be comprised of very fast random access memory (RAM) components such as, static random access memory (SRAM), asynchronous SRAM, burst SRAM, extended data output dynamic RAM (EDO

DRAM), or the like. In most cases, the high performance 40 components comprising the high speed flow cache often are relatively expensive. Thus, the high speed flow cache may comprise valuable "real estate" within a traffic management device.

The following briefly describes the various embodiments 45 to provide a basic understanding of some aspects of the invention. This brief description is not intended as an extensive overview. It is not intended to identify key or critical elements, or to delineate or otherwise narrow the scope. Its purpose is merely to present some concepts in a simplified 50 faun as a prelude to the more detailed description that is presented later.

Briefly stated, embodiments are directed towards improving the performance of network traffic management devices by optimizing the management of hot connection flows. In at 55 least one of the various embodiments, a packet traffic management device ("PTMD") may employ a data flow segment ("DFS") component and control segment ("CS") component. In at least one of the various embodiments, the CS may perform high-level control functions and per-flow policy 60 enforcement for connection flows maintained at the DFS , while the DFS may perform statistics gathering, per-packet policy enforcement (e.g., packet address translations), or the like, on connection flows maintained at the DFS.

The CS may be utilized to generate flow control data for 65 connection flows that may be offloaded to the DFS based on connection flow requests received at the packet traffic man-

4

agement device. In one embodiment, the CS may receive a new connection flow request, such as a SYN packet, sent by a client device. The CS may generate and cache a connection flow identifier for the connection flow request. In at least one of the various embodiments, the DFS may include high-speed flow caches and other high-speed components. In at least one of the various embodiments, the high-speed flow cache may be enabled to store a defined amount of flow control data that may limit the number of connection flows that may be offloaded to the DFS for handling. In at least one of the various embodiments, making efficient use of the high speed flow cache capacity may be improved by maximizing the number of hot connection flows and minimizing the number of malicious and/or in-operative connections flows (e.g., non-genuine flows) that may be have flow control data stored in the high-speed flow cache.

In at least one of the various embodiments, if a new network connection flow may be received it may be forwarded to a control segment (CS). In at least one of the various embodiments, the CS may generate the flow control data for the new network connection flow. In one embodiment, if the CS determines that the new network connection flow should be offloaded to the DFS, the CS may send a control message that may include the flow control data to the DFS. In at least one of the various embodiments, the DFS may store the received flow control data in the high-speed flow cache that may correspond to the DFS.

In at least one of the various embodiments, the CS may receive connection flows that may be evicted from the DFS. In at least one of the various embodiments, if the evicted connection may remain valid and/or active the CS may begin handling the network packets for the transferred connection flow (e.g., the CS may take over the packet level control in addition to providing the flow level control and policy enforcement).

In at least one of the various embodiments, in conjunction with managing the connection flows the CS may analyze flow statistics and application to identify hot connection flows. In at least one of the various embodiments, if hot connection flows may be identified, the CS may determine if any should be handled by the DFS for improved performance.

In at least one of the various embodiments, offloading a connection flow to the DFS for handling enables the DFS to manage packet translation using flow control data that may have generated by the CS. In at least one of the various embodiments, connection flows offloaded to the

DFS may benefit from performance improvements that arising from the high-performance hardware that may comprise the DFS. In at least one of the various embodiments, storing the flow control data for connection flows in the high-speed flow cache that may correspond to the DFS may occur if the connection flow may be offloaded to the DFS for handling.

Illustrative Operating Environment

FIG. 1 shows components of one embodiment of an environment in which the invention may be practiced. Not all of the components may be required to practice the invention, and variations in the arrangement and type of the components may be made without departing from the spirit or scope of the invention.

As shown, system 100 of FIG. 1 includes local area networks ("LANs")/wide area networks ("WANs")-(network) 108, wireless network 107, client devices 102-105, packet traffic management device ("PTMD") 109, and server devices 110-111. Network 108 is in communication with and enables communication between client devices 102-105, wireless network 107, and PTMD 109. Carrier network 107

further enables communication with wireless devices, such as client devices 103-105. PTMD 109 is in communication with network 108 and server devices 110-111.

One embodiment of client devices **102-105** is described in more detail below in conjunction with FIG. **2**. In one embodiment, at least some of client devices **102-105** may operate over a wired and/or a wireless network, such as networks **107** and/or **108**. Generally, client devices **102-105** may include virtually any computing device capable of communicating over a network to send and receive information, including instant messages, performing various online activities, or the like. It should be recognized that more or less client devices may be included within a system such as described herein, and embodiments are therefore not constrained by the number or type of client devices employed.

Devices that may operate as client device 102 may include devices that typically connect using a wired or wireless communications medium, such as personal computers, servers, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PCs, or the like. In 20 some embodiments, client devices 102-105 may include virtually any portable computing device capable of connecting to another computing device and receiving information, such as laptop computer 103, smart phone 104, tablet computer **105**, or the like. However, portable computer devices are not 25 so limited and may also include other portable devices, such as cellular telephones, display pagers, radio frequency ("RF") devices, infrared ("IR") devices, Personal Digital Assistants ("PDAs"), handheld computers, wearable computers, integrated devices combining one or more of the preceding devices, and the like. As such, client devices 102-105 typically range widely in terms of capabilities and features. Moreover, client devices 102-105 may provide access to various computing applications, including a browser, or other webbased applications.

A web-enabled client device may include a browser application that is configured to receive and to send web pages, web-based messages, and the like. The browser application may be configured to receive and display graphics, text, multimedia, and the like, employing virtually any web-based 40 language, including a wireless application protocol messages ("WAP"), and the like. In one embodiment, the browser application is enabled to employ Handheld Device Markup Language ("HDML"), Wireless Markup Language ("WML"), WMLScript, JavaScript, Standard Generalized Markup Lan- 45 guage ("SGML"), HyperText Markup Language ("HTML"), eXtensible Markup Language ("XML"), and the like, to display and send a message. In one embodiment, a user of the client device may employ the browser application to perform various activities over a network (online). However, another 50 application may also be used to perform various online activities.

Client devices 102-105 also may include at least one other client application that is configured to receive and/or send data between another computing device. The client application may include a capability to send and/or receive content, or the like. The client application may further provide information that identifies itself, including a type, capability, name, or the like. In one embodiment, client devices 102-105 may uniquely identify themselves through any of a variety of mechanisms, including a phone number, Mobile Identification Number ("MIN"), an electronic serial number ("ESN"), or other mobile device identifier. The information may also indicate a content format that the mobile device is enabled to employ. Such information may be provided in a network 65 packet, or the like, sent between other client devices, PTMD 109, server devices 110-111, or other computing devices.

6

Client devices 102-105 may further be configured to include a client application that enables an end-user to log into an end-user account that may be managed by another computing device, such as server devices 110-111, or the like. Such end-user accounts, in one non-limiting example, may be configured to enable the end-user to manage one or more online activities, including in one non-limiting example, search activities, social networking activities, browse various websites, communicate with other users, participate in gaming, interact with various applications, or the like. However, participation in online activities may also be performed without logging into the end-user account.

Wireless carrier network 107 is configured to couple client devices 103-105 and its components with network 108. Wireless carrier network 107 may include any of a variety of wireless sub-networks that may further overlay stand-alone ad-hoc networks, and the like, to provide an infrastructure-oriented connection for client devices 102-105. Such subnetworks may include mesh networks, Wireless LAN ("WLAN") networks, cellular networks, and the like. In one embodiment, the system may include more than one wireless network.

Wireless carrier network 107 may further include an autonomous system of terminals, gateways, routers, and the like connected by wireless radio links, and the like. These connectors may be configured to move freely and randomly and organize themselves arbitrarily, such that the topology of wireless carrier network 107 may change rapidly.

Wireless carrier network 107 may further employ a plurality of access technologies including 2nd (2G), 3rd (3G), 4th (4G) 5<sup>th</sup> (5G) generation radio access for cellular systems, WLAN, Wireless Router ("WR") mesh, and the like. Access technologies such as 2G, 3G, 4G, 5G, and future access networks may enable wide area coverage for mobile devices, such as client devices 103-105 with various degrees of mobility. In one non-limiting example, carrier network 107 may enable a radio connection through a radio network access such as Global System for Mobil communication ("GSM"), General Packet Radio Services ("GPRS"), Enhanced Data GSM Environment ("EDGE"), code division multiple access ("CDMA"), time division multiple access ("TDMA"), Wideband Code Division Multiple Access ("WCDMA"), High Speed Downlink Packet Access ("HSDPA"), Long Term Evolution ("LTE"), and the like. In essence, carrier network 107 may include virtually any wireless communication mechanism by which information may travel between client devices 103-105 and another computing device, network, and the like.

Network 108 is configured to couple network devices with other computing devices, including, server devices 110-111 through PTMD 109, client device 102, and client devices 103-105 through wireless carrier network 107. Network 108 is enabled to employ any form of computer readable media for communicating information from one electronic device to another. Also, network 108 can include the Internet in addition to LANs, WANs, direct connections, such as through a universal serial bus ("USB") port, other forms of computer readable media, or any combination thereof. On an interconnected set of LANs, including those based on differing architectures and protocols, a router acts as a link between LANs, enabling messages to be sent from one to another. In addition, communication links within LANs typically include twisted wire pair or coaxial cable, while communication links between networks may utilize analog telephone lines, full or fractional dedicated digital lines including T1, T2, T3, and T4, and/or other carrier mechanisms including, for example, E-carriers, Integrated Services Digital Networks ("ISDNs"),

Digital Subscriber Lines ("DSLs"), wireless links including satellite links, or other communications links known to those skilled in the art. Moreover, communication links may further employ any of a variety of digital signaling technologies, including without limit, for example, DS-0, DS-1, DS-2, 5 DS-3, DS-4, OC-3, OC-12, OC-48, or the like. Furthermore, remote computers and other related electronic devices could be remotely connected to either LANs or WANs via a modem and temporary telephone link. In one embodiment, network 108 may be configured to transport information of an Internet 10 Protocol ("IP"). In essence, network 108 includes any communication method by which information may travel between computing devices.

Additionally, communication media typically embodies computer readable instructions, data structures, program 15 modules, or other transport mechanism and includes any information delivery media. By way of example, communication media includes wired media such as twisted pair, coaxial cable, fiber optics, wave guides, and other wired media and wireless media such as acoustic, RF, infrared, and 20 other wireless media.

One embodiment of PTMD 109 is described in more detail below in conjunction with FIG. 3. Briefly, however, PTMD 109 may include virtually any network device capable of managing network traffic between client devices 102-105 and 25 server devices 110-111. Such devices include, for example, routers, proxies, firewalls, load balancers, cache devices, devices that perform network address translation, or the like, or any combination thereof. PTMD 109 may perform the operations of routing, translating, switching packets, or the 30 like. In one embodiment, PTMD 109 may inspect incoming network packets, and may perform an address translation, port translation, a packet sequence translation, and the like, and route the network packets based, at least in part, on the packet inspection. In some embodiments, PTMD 109 may 35 perform load balancing operations to determine a server device to direct a request. Such load balancing operations may be based on network traffic, network topology, capacity of a server, content requested, or a host of other traffic distribution mechanisms.

PTMD 109 may include a control segment and a separate data flow segment. The control segment may include software-optimized operations that perform high-level control functions and per-flow policy enforcement for packet traffic management. In at least one of the various embodiments, the 45 control segment may be configured to manage connection flows maintained at the data flow segment. In one embodiments, the control segment may provide instructions, such as, for example, a packet translation instruction, to the data flow segment to enable the data flow segment to route received 50 packets to a server device, such as server device 110-111. The data flow segment may include hardware-optimized operations that perform statistics gathering, per-packet policy enforcement (e.g., packet address translations), high-speed flow caches, or the like, on connection flows maintained at 55 DFS between client devices, such as client devices 102-105, and server devices, such as server devices 110-111.

Server devices 110-111 may include virtually any network device that may operate as a website server. However, server devices 110-111 are not limited to website servers, and may also operate as messaging server, a File Transfer Protocol (FTP) server, a database server, content server, or the like. Additionally, each of server devices 110-111 may be configured to perform a different operation. Devices that may operate as server devices 110-111 include various network 65 devices, including, but not limited to personal computers, desktop computers, multiprocessor systems, microprocessor-

8

based or programmable consumer electronics, network PCs, server devices, network appliances, and the like.

Although FIG. 1 illustrates server devices 110-111 as single computing devices, the invention is not so limited. For example, one or more functions of each of server devices 110-111 may be distributed across one or more distinct network devices. Moreover, server devices 110-111 are not limited to a particular configuration. Thus, in one embodiment, server devices 110-111 may contain a plurality of network devices that operate using a master/slave approach, where one of the plurality of network devices of server devices 110-111 operate to manage and/or otherwise coordinate operations of the other network devices. In other embodiments, the server devices 110-111 may operate as a plurality of network devices within a cluster architecture, a peer-to-peer architecture, and/or even within a cloud architecture. Thus, the invention is not to be construed as being limited to a single environment, and other configurations, and architectures are also envisaged.

#### Illustrative Client Device

FIG. 2 shows one embodiment of client device 200 that may be included in a system implementing embodiments of the invention. Client device 200 may include many more or less components than those shown in FIG. 2. However, the components shown are sufficient to disclose an illustrative embodiment for practicing the present invention. Client device 200 may represent, for example, one embodiment of at least one of client devices 102-105 of FIG. 1.

As shown in the figure, client device 200 includes a processor 202 in communication with memory 226 via a bus 234. Client device 200 also includes a power supply 228, one or more network interfaces 236, an audio interface 238, a display 240, a keypad 242, and an input/output interface 248.

Power supply 228 provides power to client device 200. A rechargeable or non-rechargeable battery may be used to provide power. The power may also be provided by an external power source, such as an AC adapter or a powered docking cradle that supplements and/or recharges a battery.

Client device 200 may optionally communicate with a base station (not shown), or directly with another computing device. Network interface 236 includes circuitry for coupling client device 200 to one or more networks, and is constructed for use with one or more communication protocols and technologies including, but not limited to, global system for mobile communication ("GSM"), code division multiple access ("CDMA"), time division multiple access ("TDMA"), High Speed Downlink Packet Access ("HSDPA"), Long Term Evolution ("LTE"), user datagram protocol ("UDP"), transmission control protocol/Internet protocol ("TCP/IP"), short message service ("SMS"), general packet radio service ("GPRS"), WAP, ultra wide band ("UWB"), IEEE 802.16 Worldwide Interoperability for Microwave Access ("WiMax"), session initiated protocol/real-time transport protocol ("SIP/RTP"), or any of a variety of other wireless communication protocols. Network interface 236 is sometimes known as a transceiver, transceiving device, or network interface card ("NIC")

Audio interface 238 is arranged to produce and receive audio signals such as the sound of a human voice. For example, audio interface 238 may be coupled to a speaker and microphone (not shown) to enable telecommunication with others and/or generate an audio acknowledgement for some action.

Display 240 may be a liquid crystal display ("LCD"), gas plasma, light emitting diode ("LED"), or any other type of display used with a computing device. Display 240 may also

include a touch sensitive screen arranged to receive input from an object such as a stylus or a digit from a human hand.

Keypad 242 may comprise any input device arranged to receive input from a user. For example, keypad 242 may include a push button numeric dial, or a keyboard. Keypad 242 may also include command buttons that are associated with selecting and sending images.

Client device 200 also comprises input/output interface 248 for communicating with external devices, such as a head-set, or other input or output devices not shown in FIG. 2. Input/output interface 248 can utilize one or more communication technologies, such as USB, infrared, Bluetooth<sup>TM</sup>, or the like

Client device 200 may also include a GPS transceiver (not shown) to determine the physical coordinates of client device 200 on the surface of the Earth. A GPS transceiver typically outputs a location as latitude and longitude values. However, the GPS transceiver can also employ other geo-positioning mechanisms, including, but not limited to, triangulation, 20 assisted GPS ("AGPS"), Enhanced Observed Time Difference ("E-OTD"), Cell Identifier ("CI"), Service Area Identifier ("SAI"), Enhanced Timing Advance ("ETA"), Base Station Subsystem ("BSS"), or the like, to further determine the physical location of client device 200 on the surface of the 25 Earth. It is understood that under different conditions, a GPS transceiver can determine a physical location within millimeters for client device 200; and in other cases, the determined physical location may be less precise, such as within a meter or significantly greater distances. In one embodiment, however, mobile device 200 may through other components, provide other information that may be employed to determine a physical location of the device, including for example, a Media Access Control ("MAC") address, IP address, or the

Memory 226 includes a Random Access Memory ("RAM") 204, a Read-only Memory ("ROM") 222, and other storage means. Mass memory 226 illustrates an example of computer readable storage media (devices) for storage of information such as computer readable instructions, data 40 structures, program modules or other data. Mass memory 226 stores a basic input/output system ("BIOS") 224 for controlling low-level operation of client device 200. The mass memory also stores an operating system 206 for controlling the operation of client device 200. It will be appreciated that 45 this component may include a general-purpose operating system such as a version of UNIX, or LINUX<sup>TM</sup>, or a specialized client communication operating system such as Windows Mobile™, or the Symbian® operating system. The operating system may include, or interface with a Java virtual machine 50 module that enables control of hardware components and/or operating system operations via Java application programs.

Mass memory 226 further includes one or more data storage 208, which can be utilized by client device 200 to store, among other things, applications 214 and/or other data. For 55 example, data storage 208 may also be employed to store information that describes various capabilities of client device 200. The information may then be provided to another device based on any of a variety of events, including being sent as part of a header during a communication, sent upon 60 request, or the like. Data storage 208 may also be employed to store social networking information including address books, buddy lists, aliases, user profile information, or the like. Further, data storage 208 may also store message, we page content, or any of a variety of user generated content. At least a 65 portion of the information may also be stored on another component of network device 200, including, but not limited

10

to processor readable storage device 230, a disk drive or other computer readable storage medias (not shown) within client device 200.

Processor readable storage device 230 may include volatile, nonvolatile, removable, and non-removable media implemented in any method or technology for storage of information, such as computer- or processor-readable instructions, data structures, program modules, or other data. Examples of computer readable storage media include RAM, ROM, Electrically Erasable Programmable Read-only Memory ("EEPROM"), flash memory or other memory technology, Compact Disc Read-only Memory ("CD-ROM"), digital versatile disks ("DVD") or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other physical medium which can be used to store the desired information and which can be accessed by a computing device. Processor readable storage device 230 may also be referred to herein as computer readable storage media.

Applications 214 may include computer executable instructions which, when executed by client device 200, transmit, receive, and/or otherwise process network data. Network data may include, but is not limited to, messages (e.g., SMS, Multimedia Message Service ("MMS"), instant message ("IM"), email, and/or other messages), audio, video, and enable telecommunication with another user of another client device. Applications 214 may include, for example, browser 218. Applications 214 may include other applications, which may include, but are not limited to, calendars, search programs, email clients, IM applications, SMS applications, voice over Internet Protocol ("VOIP") applications, contact managers, task managers, transcoders, database programs, word processing programs, security applications, spreadsheet programs, games, search programs, and so forth.

Browser 218 may include virtually any application configured to receive and display graphics, text, multimedia, and the like, employing virtually any web based language. In one embodiment, the browser application is enabled to employ HDML, WML, WMLScript, JavaScript, SGML, HTML, XML, and the like, to display and send a message. However, any of a variety of other web-based programming languages may be employed. In one embodiment, browser 218 may enable a user of client device 200 to communicate with another network device, such as PTMD 109 and/or with server devices 110-111.

#### Illustrative Network Device

FIG. 3 shows one embodiment of a network device 300, according to one embodiment of the invention. Network device 300 may include many more or less components than those shown. The components shown, however, are sufficient to disclose an illustrative embodiment for practicing the invention. Network device 300 may be configured to operate as a server, client, peer, a host, or any other device. Network device 300 may represent, for example PTMD 109 of FIG. 1, server devices 110-111 of FIG. 1, and/or other network devices.

Network device 300 includes processor 302, processor readable storage device 328, network interface unit 330, an input/output interface 332, hard disk drive 334, video display adapter 336, data flow segment ("DFS") 338 and a mass memory, all in communication with each other via bus 326. The mass memory generally includes RAM 304, ROM 322 and one or more permanent mass storage devices, such as hard disk drive 334, tape drive, optical drive, and/or floppy disk drive. The mass memory stores operating system 306 for controlling the operation of network device 300. Any general-purpose operating system may be employed. Basic input/

output system ("BIOS") 324 is also provided for controlling the low-level operation of network device 300. As illustrated in FIG. 3, network device 300 also can communicate with the Internet, or some other communications network, via network interface unit 330, which is constructed for use with various communication protocols including the TCP/IP protocol. Network interface unit 330 is sometimes known as a transceiver, transceiving device, or network interface card ("NIC").

Network device **300** also comprises input/output interface **332** for communicating with external devices, such as a keyboard, or other input or output devices not shown in FIG. **3**. Input/output interface **332** can utilize one or more communication technologies, such as USB, infrared, Bluetooth<sup>TM</sup>, or the like.

The mass memory as described above illustrates another type of computer readable media, namely computer readable storage media and/or processor readable storage media, including processor readable storage device 328. Processor readable storage device 328 may include volatile, nonvolatile, removable, and non-removable media implemented in any method or technology for storage of information, such as computer readable instructions, data structures, program modules, or other data. Examples of processor readable storage media include RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other media which can be used to store the desired information and which can be accessed by a computing device.

Data storage 308 may include a database, text, spreadsheet, folder, file, or the like, that may be configured to maintain and store user account identifiers, user profiles, email addresses, 35 IM addresses, and/or other network addresses; or the like. Data stores 308 may further include program code, data, algorithms, and the like, for use by a processor, such as central processing unit 302 to execute and perform actions. In one embodiment, at least some of data store 308 might also be 40 stored on another component of network device 300, including, but not limited to processor-readable storage device 328, hard disk drive 334, or the like.

The mass memory may also stores program code and data. One or more applications **314** may be loaded into mass 45 memory and run on operating system **306**. Examples of application programs may include transcoders, schedulers, calendars, database programs, word processing programs, Hypertext Transfer Protocol ("HTTP") programs, customizable user interface programs, IPSec applications, encryption programs, security programs, SMS message servers, IM message servers, email servers, account managers, and so forth. Web server **316** and control segment ("CS") **318** may also be included as application programs within applications **314**.

Web server **316** represent any of a variety of services that 55 are configured to provide content, including messages, over a network to another computing device. Thus, web server **316** includes, for example, a web server, a File Transfer Protocol ("FTP") server, a database server, a content server, or the like. Web server **316** may provide the content including messages 60 over the network using any of a variety of formats including, but not limited to WAP, HDML, WML, SGML, HTML, XML, Compact HTML ("cHTML"), Extensible HTML ("xHTML"), or the like. Web server **316** may also be configured to enable a user of a client device, such as client devices 65 **102-105** of FIG. **1**, to browse websites, upload user data, or the like.

12

Network device 300 may also include DFS 338 for maintaining connection flows between client devices, such as client devices 102-105 of FIG. 1, and server devices, such as server devices 110-111 of FIG. 1. In one embodiment, DFS 338 may include hardware-optimized operations for packet traffic management, such as repetitive operations associated with packet traffic management. For example, DFS 338 may perform statistics gathering, per-packet policy enforcement (e.g., packet address translations), or the like, on connection flows maintained at DFS 338. In some embodiments, DFS 338 may route, switch, forward, direct and/or otherwise handle packets based on rules for a particular connection flow signature (e.g., a 5 tuple of a received packet). Thus, DFS 338 may include capabilities and perform tasks such as that of a router, a switch, a routing switch, or the like. In some embodiments, the rules for a particular connection flow signature may be based on instructions received from CS 318. In one embodiment, DFS 338 may store the instructions received from CS 318 in a local memory as a table or some other data structure. In some other embodiments, DFS 338 may also store a flow state table to indicate a state of current connection flows maintained at DFS 338. In at least one of the various embodiments, components of DFS 338 may comprise and/or work in combination to provide high-speed flow caches for optimizing packet traffic management.

In some embodiments, DFS 338 may provide connection flow status updates to CS 318. In one embodiment, a connection flow status update may include a status of the connection flow, a current state of the connection flow, other statistical information regarding the connection flow, or the like. The connection flow update may also include an identifier that corresponds to the connection flow. The identifier may be generated and provided by CS 318 when a connection flow is established at DFS 338. In some embodiments, the connection flow update may be a connection flow delete update provided to CS 318 after the connection flow is terminated at DFS 338. The connection flow status update and/or the connection flow delete update may be provided to CS 318 periodically, at predefined time intervals, or the like. In some embodiments, DFS 338 may stagger a time when a plurality of connection flow status updates are provided to CS.

In some other embodiments, DFS 338 may include a plurality of data flow segments. In one non-limiting example, a first data flow segment within DFS 338 may forward packets received from a client device to a server device, while a second data flow segment within DFS 338 may forward and/or route packets received from a server device to a client device. In at least one of the various embodiments, DFS 338 may also be implemented in software.

CS 318 may include a control segment that may include software-optimized operations to perform high-level control functions and per-flow policy enforcement for packet traffic management. CS 318 may be configured to manage connection flows maintained at DFS 338. In one embodiments, CS 318 may provide instructions, such as, for example, a packet address translation instructions, to DFS 338 to enable DFS 338 to forward received packets to a server device, such as server device 110-111 of FIG. 1. In some other embodiments, CS 318 may forward and/or route packets between a client device and a server device independent of DFS 338.

In at least one of the various embodiments, CS 318 may include a plurality of control segments. In some embodiments, a plurality of control segments may access and/or manage connection flows at a single data flow segments and/or a plurality of data flow segments. In some other embodiments, CS 318 may include an internal data flow segment. In one such embodiment, the internal data flow segment of CS

318 may be distributed and/or separate from CS 318. For example, in one embodiment, CS 318 may be employed in software, while the internal data flow segment may be employed in hardware. In some other embodiments, CS 318 may identify if connection flows are split between different 5 data flow segments and/or between a DFS 338 and CS 318. In at least one embodiment, CS 318 may also be implemented in hardware.

In at least one of the various embodiments, CS 318 may be configured to generate an identifier for each connection flow established at DFS 338. In some embodiments, CS 318 may utilize a sequence number of a SYN to generate an identifier for a corresponding connection flow.

In one embodiment, the identifier may be based on a hash of the sequence number. In another embodiment, the identifier may be based on an exclusive OR byte operation of the sequence number. CS 318 may cache the identifier at CS 318 and may provide the identifier to DFS 338. In some embodiments, CS 318 may cache an identifier for each connection flow it establishes at DFS 338.

FIG. 4A illustrates a system diagram generally showing one embodiment of a system with a packet traffic management device disposed between client devices and server devices. System 400A may include packet traffic management device ("PTMD") 404 disposed between client devices 25 **402-403** and server devices **410-411**. Client devices **402-403** may include Client\_1 through Client\_M, which may include one or more client devices, such as client devices 200 of FIG. 2. Server devices 410-411 may include Server 1 through Server\_N, which may include one or more server devices, 30 such as server devices 110-111 of FIG. 1.

In one embodiment, PTMD 404 may be an embodiment of PTMD 109 of FIG. 1. PTMD 404 may include data flow segment ("DFS") 406 in communication with control segment ("CS") 408. In at least one of the various embodiments, 35 DFS 406 may be an embodiment of DFS 338 of FIG. 3, and CS 408 may be an embodiment of CS 318 of FIG. 3.

CS 408 may be configured to communicate with DFS 406, client devices 402-403 and/or server devices 410-411 indemay establish connection flows at DFS 406. In some embodiments, CS 408 may establish a connection flow at DFS 406 by providing instructions including flow control data to DFS 406 that enables DFS 406 to forward packets received at PTMD 404. In one embodiment, CS 408 may perform a load balanc- 45 ing operation to select a server device of server devices 410-411 to receive packets sent from a client device, such as client device 402. In some other embodiments, CS 408 may generate and cache a connection flow identifier to be provided to DFS 406 when the connection flow is established.

DFS 406 may be configured to facilitate communications between client devices 402-403 and server devices 410-411. DFS 406 may process and forward packets received at PTMD **404** based on the instructions and flow control data received from CS 408. For example, in one embodiment, DFS 406 55 utilizes the instructions and/or flow control data to forward packets received from client device 402 to server device 410 and to forward packets received from server device 410 to client device 402. In some embodiments, DFS 406 may forward predetermined packets to CS 408, such as, but not 60 limited to, new connection flow requests (e.g., associated with a SYN). In yet other embodiments, DFS 406 may notify CS 408 that a packet was received and forwarded. In one non-limiting, non-exhaustive example, DFS 406 may notify CS 408 that an ACK was received from client device 402 and 65 forwarded to server device 410. In at least one of the various embodiments, DFS 406 may also provide connection flow

14

updates and a corresponding connection flow identifier to CS 408. CS 408 may compare the corresponding connection flow identifier with the cached identifier to determine if the connection flow update is valid.

In at least one of the various embodiments, DFS 406 may send evict messages to CS 408 if connection flow are evicted from the DFS 406. In at least one of the various embodiments, DFS 406 may evict a connection flow if new flows arrive and the capacity of the DFS to handle new connection flow may be exceeded. In at least one of the various embodiments, evictions from DFS 406 may occur if the high speed flow cache for storing flow control data exhausts its ability to store the flow control data for new connection flows. In at least one of the various embodiments, evict messages sent from DFS 406to CS 408 may contain enough information to fully identify the connection flow (e.g., endpoints, ports, sequent numbers, flow state, or the like).

In at least one of the various embodiments, CS 408 may receive and route packets associated with evicted connection flows, thereby taking on some of the duties of DFS 406. In at least one of the various embodiments, some new connection flow may not be offloads to DFS 406 if CS 408 determines that the connection flows may be management on the CS or if the CS determines that more information may be required to determine if the connection flow should be offloaded to DFS

Although PTMD 404 illustrates DFS 406 and CS 408 as two partitions within a single PTMD 404, the invention is not so limited. Rather, in some embodiments, DFS 406 and CS 408 may be functional blocks in a same PTMD 404 (i.e., a same chassis/computing device). In other embodiments, DFS 406 may be implemented by one or more chassis/computing devices separate from one or more other chassis/computing devices that may be utilized to implement CS 408. In yet other embodiments, CS 408 may be a module that plugs into DFS **406**. Additionally, it is envisaged that the functionality of either DFS 406 and/or CS 408 may be separately implemented in software and/or hardware.

FIG. 4B illustrates a system diagram generally showing pendent of DFS 406, and/or any combination thereof. CS 408 40 one embodiment of a system with a packet traffic management device disposed between client devices and server devices. System 400B may include packet traffic management device ("PTMD") 404 disposed between client devices 402-403 and server devices 410-411. Client devices 402-403 may include Client 1 through Client\_M, which may include one or more client devices, such as client devices 102-105 of FIG. 1. Server devices 410-411 may include Server 1 through Server N, which may include one or more server devices, such as server devices 110-111 of FIG. 1.

> In one embodiment, PTMD 404 may be an embodiment of PTMD 404 of FIG. 4. PTMD 404 may include data flow segments ("DFS") 406-407 and control segments ("CS") 408-409. DFS 406-407 may include a plurality of data flow segments, each of which may be an embodiment of DFS 406 of FIG. 4A. CS 408-409 may include a plurality of control flow segments, each of which may be an embodiment of CS

> In some embodiments, data communicated between client devices 402-403 and server devices 410-411 may flow through one or more data flow segments 406-407. In one embodiment, data from client devices 402-403 may flow through a first DFS, such as DFS 406 and data from server devices 410-411 may flow through a second DFS, such as DFS 407.

In at least one of the various embodiments, one or more data flow segments of DFS 406-407 may communicate with one or more control segments of CS 408-409. Similarly, one

or more control segments of CS **408-409** may communicate with one or more data flow segments of DFS **406-407**. In some embodiments, each control segment of CS **408-409** may communicate (not shown) with other control segments of CS **408-409**. In other embodiments, each data flow segment of DFS **406-407** may communicate (not shown) with other data flow segments of DFS **406-407**.

Also, in at least one of the various embodiments, connection flows may be split into flow portions based on the direction of network packet travel. In at least one of the various embodiments, the network packets coming from the client may treated as a separate connection flow and the network packets coming from a server and directed towards a client may be treated as a separate connection flow. In at least one of the various embodiments, this enables optimizations based on 15 the amount of network packet traffic of a particular split connection flows. In at least one of the various embodiments, this may enable the upload and download direction portion of connection flows to be split across CS 408-409 and DFS **406-407** based on the characteristics of the upload and down- 20 load portions of the connection flows. For example, in at least one of the various embodiments, if downloading streaming video may be a very asymmetric operation having many network packets download to the client and few uploaded. In at least one of the various embodiments, the upload and 25 download portions of connection flow in the download direction may be optimized independent with one portion using the DFS and a high-speed flow cache and the other portion may be handled on the CS using lower performing (e.g., less expensive) resources.

FIG. 5 illustrates an embodiment of a sequence for establishing a connection flow and offloading the new connection flow to the data flow segment (DFS). Sequence 500 may show an embodiment using TCP/IP networking protocol but one of ordinary skill the art will appreciate that the sequence dia- 35 gram (or similar sequences) may generally apply to other networking protocols that may have other handshaking sequences as well. Also, even though sequence 500 depicts a sequence including one client, one DFS, one CS, and one application server, in at least one of the various embodiments, 40 one or more, data flow segments, control segments, clients, and servers, may be participate in handshaking and in the connection flow offloading. Also, in at least one of the various embodiments, the connection flows may be split into upload and download portions of a connection flow, with each por- 45 tion representing one direction of the connection flow.

In at least one of the various embodiments, sequence **500** begins at step **502** if a client initiates a connection with a network resource that may be managed by a PTMD, such as PTMD **109**. If client may be initiating the connection using 50 TCP/IP, a SYN packet may be sent to the PTMD.

At step **504** a SYN packet may be received at a DFS that may be part of a PTMD. In at least one of the various embodiments, at step **506**, because the DFS may determine that the incoming connection represents a new connection flow, the 55 DFS may forward the SYN packet to a CS. At step **506** a CS may examine the connection flow and may determine the appropriate flow control data for the new flow and send it to the DFS. In at least one of the various embodiments, CS may apply one or more stored rules that may be used to determine 60 the flow control data for the new network connection flow. In at least one of the various embodiments, the stored rules may implement network traffic management services such as load balancing, application access control, or the like.

In at least one of the various embodiments, at step **508** the 65 DFS may receive the flow control data from the CS and store it in a high speed flow cache. In at least one of the various

16

embodiments, the flow control data may be used by the DFS to forward the SYN packet to an appropriate server and/or network resource as directed by the flow control data that may be provided by the CS.

In at least one of the various embodiments, at step 510 a server and/or network resource may receive the SYN packet and may respond by sending a SYN-ACK packet to the DFS. In at least one of the various embodiments, at step 512 the DFS may again use the flow control data stored in the high speed flow cache to map and/or translate the SYN\_ACK from a server to the appropriate client.

In at least one of the various embodiments, at step **514** the client device that sent the initial SYN packet may receive the corresponding SYN\_ACK and subsequently may respond with an ACK packet. In at least one of the various embodiments, at step **516** the DFS, using the stored flow control data to determine the network path the to server, may forward the ACK packet to the server.

In at least one of the various embodiments, at step 518 the server may receive the ACK packet corresponding to the client device. After the ACK may have been received, the network connection flow may be in an established state. In at least one of the various embodiments, during steps 520-524, using the established network connection flow, the server may begin exchanging application data with client. In at least one of the various embodiments, at this point, for each exchange of data, the DFS may use the flow control data that may be stored in the high speed flow cache to map between the application servers and the client to route the packets on the correct path to maintain the connection flow.

#### General Operation

FIG. 6 shows a flowchart showing at least one of the various embodiments of a process for packet traffic management. In process 600, after a start block, at block 602 a network packet may be received by a DFS. In at least one of the various embodiments, the network packets may be received from network 108, and/or may have been forwarded through multiple networks, switches, routers, other PTMDs or the like.

At decision block **604**, in at least one of the various embodiments, if the received network packet may be associated with a new connection flow, control may move to block **606**. Otherwise, in at least one of the various embodiments, control may move to decision block **608**.

In at least one of the various embodiments, a DFS may examine the connection flow and compare it the flow control data that may be stored in a high-speed cache. In at least one of the various embodiments, a tuple corresponding to the network packet may be examined to determine if the network packet is part of a new connection flow. If a tuple corresponding to the incoming network packet may not be found in the high-speed flow cache the DFS may determine that the network packet may be part of a new connection flow.

At block 606, in at least one of the various embodiments, the incoming network packet that may be associated with a new connection flow may be forwarded to a CS for further processing. In at least one of the various embodiments, the incoming network packet may be sent to a CS using a command bus that may enable DFS and CS components to exchange data and messages. Next, control may move decision block 614.

At decision block 608, in at least one of the various embodiments, if flow control data may be available for the connection flow associated with network packet, control may move to block 710. Otherwise, in at least one of the various embodiments, control may move to block 612.

At block 610, in at least one of the various embodiments, the DFS may forward the network packet to its next destina-

tion based on the flow control data and/or information associated with the network packet's corresponding connection flow that may be stored in the high speed flow cache that corresponds to the DFS. Next, in at least one of the various embodiments, control may move to decision block **614**.

At block **612**, in at least one of the various embodiments, the network packet having a previously seen tuple may be stored in a buffer on the DFS until flow control data may be provided by the CS.

In at least one of the various embodiments, a received 10 network packet may be associated with a connection flow that has been previously been observed. However, in at least one of the various embodiments, if the flow control data from the CS may not be available, the DFS may store the network packets associated with the connection flow in a buffer until 15 the relevant flow control data may be received from the CS.

Also, in at least one of the various embodiments, incoming network packets associated with unknown and/or new connection flows may be forwarded to the CS for buffering, rather than buffering on the DFS, until a flow control data determination may be made by the CS.

At decision block **614**, in at least one of the various embodiments, if there may be more incoming network packets, control may loop back to block **602**. Otherwise, in at least one of the various embodiments, control may be returned to a 25 calling process.

FIG. 7 shows a flowchart of process 700, in at least one of the various embodiments, for handling new connection flows at a DFS. After a start block, at block 702 a DFS component may receive new flow control data from a CS. In at least one 30 of the various embodiments, if new flow control data may be received, the DFS may store the flow data into a high-speed flow cache.

In at least one of the various embodiments, the new flow control data may be sent to the DFS as part of a "new flow" 35 control message sent from the CS to the DFS.

At decision block **704**, in at least one of the various embodiments, if the DFS high-speed flow cache may be full, control may move to block **706**. Otherwise, in at least one of the various embodiments, control may move block **708**.

In at least one of the various embodiments, the high-speed flow cache may be implemented as a hash such that the a hash key may be generated for each new connection flow based on properties of the connection flow such as the tuple, CS generated connection identifier, SYN cookie, or the like. In at 45 least one of the various embodiments, if the range (number of unique values) of the hash key may be more than the number of slots in the high speed flow cache, the hash key may be truncated so the number of hash key value possibilities may be equal or similar to the number of slots in the high-speed 50 flow cache. In at least one of the various embodiments, truncation of the hash key may increase the number of hash key collisions. If, in at least one of the various embodiments, a new connection flow hash key may cause hash key collision, the connection flow currently in the cache may get evicted 55 (e.g., its flow control data is removed from the high speed cache and the responsibility for managing the flow may be transferred to the CS) to make room for the new connection

At block **706**, in at least one of the various embodiments, to 60 make room for the new flow control data received from the CS, flow control data for a different, previously cached connection flow may be removed (e.g., evicted) from the DFS high-speed flow cache. In at least one of the various embodiments, the DFS may send the CS a control message indicating 65 that a connection flow may have been evicted from the DFS requiring the associated flow control data to be removed from

18

the DFS high-speed flow cache. In at least one of the various embodiments, the eviction message may include information, such as, number of packets sent or received over this network flow, age of the network flow, tuple information, or the like. In at least one of the various embodiments, the control message sent to the CS may contain enough information to enable the CS to identify the network flow that may be evicted from the DFS

At block 708, in at least one of the various embodiments, the flow control data associated with the new connection flow may be stored in the DFS high-speed flow cache. In at least one of the various embodiments, flow control data may be stored in one or more components of the DFS that may operate singly or in combination as a high-speed flow cache.

At block 710, in at least one of the various embodiments, the DFS may begin processing received network packets associated with known connection flows using the flow control data that may be associated with the connection flow and stored in the high-speed flow cache.

FIG. **8** shows a flowchart of process **800**, in at least one of the various embodiments, for handling eviction (EVICT) messages at a CS. After a start block, at block **802**, the CS may receive an EVICT message from a DFS.

At decision block **804**, in at least one of the various embodiments, if the eviction message corresponds to a closed and/or terminated connection flow control may move to block **806**. Otherwise, in at least one of the various embodiments, control may move to block **808**.

At block **806**, in at least one of the various embodiments, the closed and/or terminated flow and associated flow control data may be discarded.

At block **808**, in at least one of the various embodiments, the responsibility for managing the evicted connection flow may be transferred to the CS. In at least one of the various embodiments, network packets received over the transferred connection flow may be handled by the CS.

In at least one of the various embodiments, the CS may store the flow control data for the evicted connection flow in a local flow cache. In at least one of the various embodiments, the flow cache in the CS may be arranged to include at least the same information that may be stored regarding connection flows using the high speed flow cache on the DFS.

At decision block 810, in at least one of the various embodiments, if there may be more eviction messages to process, control may loop back to block 802. Otherwise, in at least one of the various embodiments, control may be returned to a calling process.

In at least one of the various embodiments, depending on the circumstances, a connection flow may be handled on one or more DFSs, on one or more CSs, or partially on one or more CSs and partially on one or more DFSs. In at least one of the various embodiments, if a connection flow may be being handled by the CS it may not receive a new flow network message from the DFS. Likewise, if a DFS may be handling a connection flow it may not send a new flow network message to the CS component if the DFS can associate the incoming network traffic with a known connection flow. However, in at least one of the various embodiments, the CS may analyze each connection flow to determine the connection flows may be evicted from the DFS.

FIG. 9 shows a flowchart for process 900 that in at least one of the various embodiments determines if connection flows may be candidates for off-loading to the DFS for handling. After a start block, at block 902, in at least one of the various embodiments, the CS may receive a network packet associated with a connection flow that may be managed by the CS.

In at least one of the various embodiments, network packets received by the CS may be associated connection flows that may have their packet level processing and management processing handled on the CS rather the DFS. In at least one of the various embodiments, as the CS handles the received packets at least in accordance with the stored flow control data it may perform additional action to identify hot connection flows

At block **904**, in at least one of the various embodiments, the CS may receive a flow status update (FSU) from a DFS. In at least one of the various embodiments, the FSU may be received asynchronously with respect to the network packets that may be received by the CS. In at least one of the various embodiments, if a FSU may be not be available control may move to block **906**.

At block 906, in at least one of the various embodiments, the CS may update the statistics being maintained for the connection flows. In at least one of the various embodiments, statistics may be tracked for the connection flows being managed by the CS directly as well as the connection flows that may be managed by the DFS (e.g., off-loaded connection flows).

In at least one of the various embodiments, the updating of connection flow metrics may use a combination of information from one or more FSUs and metrics that may be collected on the CS, including, bit-rate, data sent over a time interval, data received over a time interval, or the like. In at least one of the various embodiments, the connection flow metrics collected may be based, low level network information derived from L1-L4 as well as higher level network information derived from L5-L7 (as per the Open Systems Interconnection (OSI) model).

At block **908**, in at least one of the various embodiments, the CS may analyze the collected connection flow statistics 35 and may apply relevant rules to identify hot connection flows. (See, FIGS. **10** and **11**.) In at least one of the various embodiments, the CS may employ at least one connection flow metric to determine each hot connection flow out of the plurality of managed connection flows

In at least one of the various embodiments, rules may be defined that declare that one or more specific sources, endpoints, data types, or the like may indicated to be hot connection flows. Or, in at least one of the various embodiments, rules may be defined to adjust the priority of certain connection flows based on flow patterns, sources, endpoints, data types, or like.

In at least one of the various embodiments, generally the same type of flow control policies rulemaking may be extended to influence the identification and determination of 50 how connection flows may be designated as hot connection flows.

At decision block 910, if connection flows may be identified for moving from the CS to the DFS and/or from the DFS to the CS for handling, control may move block 912. Otherwise, in at least one of the various embodiments, control may move to decision block 914.

In at least one of the various embodiments, the CS may employ at least one connection flow metric to determine each hot connection flow in the plurality of managed connection 60 flows.

At block 912, in at least one of the various embodiments, if connection flows may be identified for moving, the CS may generate the relevant commands and/or messages to and send to the appropriate CS and/or DFS for handling. In at least one of the various embodiments, some connection flows may be moved from a DFS to the CS for handling. In at least one of

20

the various embodiments, the DFS may be employed to handle each determined hot connection flow.

Also, in at least one of the various embodiments, some of the connection flows that may have been identified as hot connection flows may be moved and/or off-loaded to a DFS for handling to benefit from at least higher performance/ processing speeds the may be associated with a DFS. Next, in at least one of the various embodiments, control may move to decision block 914.

At decision block 914, in at least one of the various embodiments, if there may be more network packets available control may loop back to block 902. Otherwise, in at least one of the various embodiments, control may be returned to a calling process.

FIGS. 10 and 11 describe various embodiments for identifying if a connection flow may be a hot connection flow. One of ordinary skill in the art will appreciate that the techniques, parameters, and thresholds used to identify "hot connection flows" may vary depending on the applications being managed by a PTMD and the goals and priorities of the operators and users of the PTMD at a particular time. In at least one of the various embodiments, generally, the criteria for identifying a hot connection flow may be defined based on the application and user goals and if connection flow properties meet the criteria, a connection flow may be deemed a hot connection flow.

In at least one of the various embodiments, as part of determining if a connection flow may be a candidate for offloading to the DFS for handling (e.g., hot connection flow), the content of received network packets may be examined. In at least one of the various embodiments, the network packets may be examined to identify data patterns and meta-data that may indicate that the connection flow may be a hot connection flow that may be a good candidate for offloading to the DFS component.

In at least one of the various embodiments, if examining the network packets, the CS may identify application level protocol data, messages, or meta-data for determining if the associated connection flow may be a hot connection flow. For example, if a CS may identify that a connection flow may be using HTTP, the CS may examine HTTP headers such as, Content-Type, Content-Length, Cache-Control, or the like, as part of determining if a connection flow may be a hot connection flow.

In at least one of the various embodiments, if a network packet may be determined to be a first packet of a HTTP response, a content length value provided by the server sending the HTTP response may be available. In at least one of the various embodiments, the HTTP content length value may indicate the number of network packets that may be likely to be used to transmit the complete HTTP response from the server. For example, in at least one of the various embodiments, if the content length value may indicate that the response may use a single network packet, the associated connection flow may not be a candidate for offloading to the DFS because additional packets may not be expected for this response. On the other hand, in at least one of the various embodiments, if the content length value indicates that more network packets may be on the way for the same response, the connection flow may be determined to be a candidate for offloading to the DFS component. In at least one of the various embodiments, the content length value may correlate to the likelihood of offloading a connection flow to a DFS (e.g., an increase in the content length value leads to an increase in the chance of offloading the connection flow to the DFS).

In at least one of the various embodiments, in some cases, the operating characteristics of a connection flow may have

significant variance. For example, in at least one of the various embodiments, the bit-rate for a connection may be prone to spikes if the content/communication may be uneven. Thus, in at least one of the various embodiments, a connection flow once determined to be a good offload candidate (leading to likely offloading to the DFS) may soon be determined to be a poor offload candidate (leading to likely removal from the DFS) depending on the immediate condition and/or characteristics of the underlying communication session.

In at least one of the various embodiments, a connection flow may repeatedly cycled back and forth from being handled on the DFS to being handled on CS, or back again. In at least one of the various embodiments, the cycling may occur based on at least the variance of the operating characteristics of the connection flow. In at least one of the various embodiments, this at least enables the performance of the connection flow and the usage of the DFS to be continually optimized to take advantage of the variance in the connection flow operation.

For example, in at least one of the various embodiments, as the network traffic over the connection flow slows, the connection flow may be moved to the CS for handling. Likewise, in at least one of the various embodiments, as the network traffic over the connection flow increases the connection flow 25 may be moved to the DFS for handling.

In at least one of the various embodiments, the cycling of the connection flow between the CS and the DFS components may occur one or more times during a communication session. Also, in at least one of the various embodiments, the cycling operations performed by the CS and the DFS components may be seamless and unseen/opaque to both ends of the communication session.

FIG. 10 shows a flowchart for at least one of the various embodiments of process 1000 for identifying hot connection flows. After a start block, at decision block 1002, in at least one of the various embodiments, if the number of connection flows being handled on the PTMD may less than the capacity of the DFS control may be returned to the calling process. 40 Otherwise, in at least one of the various embodiments, control may move to block 1004. In at least one of the various embodiments, if the high speed flow cache on the DFS has unused capacity both hot connection flows and "normal" connection flows may be processed on the DFS.

At block 1004, in at least one of the various embodiments, connection flows may be sorted in rank order based on the amount of data traffic passed through, exchanged, or communicated through the connection flow in a given time interval.

In at least one of the various embodiments, well known data structures and sorting algorithms may be employed to generate a tabular data structure wherein the connection flows may be logically order from based on amount of the data traffic passing through the connection flow over a time inter-

At block 1006, in at least one of the various embodiments, hot flow candidates may be determined and identified based on the top N flows based on the rank order.

In at least one of the various embodiments, the rules associated with determining/defining hot connection flows may include parameters such as "N" (e.g., how many of the top connection flows may be designated as hot connection flows). In at least one of the various embodiments, "N" may be based on a formula that may include additional parameters including having different values based on the type of connection flow.

22

Next, control may be returned to a calling process.

FIG. 11 shows a flowchart for at least one of the various embodiments of process 1100 for identifying hot connection flows. After a start block, at decision block 1102, in at least one of the various embodiments, if the number of connection flows being handled on the PTMD may be less than the capacity of the DFS, control may be returned to the calling process. Otherwise, in at least one of the various embodiments, control may move to block 1104. In at least one of the various embodiments, if the high speed flow cache on the DFS has unused capacity both hot connection flows and "normal" connection flows may be processed on the DFS.

At block 1104, in at least one of the various embodiments, the median bit-rate of connection flows being handled on the 15 CS may be determined for use in predicting a maximum number of connection flows that may be processed by the CS. For example, in at least one of the various embodiments, if the median bit-rate of connection flows currently being handled on the CS may be 1 million bits per second and the total 20 bandwidth of the CS for handling connection flows may be 2000 million bits per second, the maximum number of connection flows that may be processed may be estimated as 2000 connection flows (2000 million bits/sec/1 million bits/sec).

At block 1106, in at least one of the various embodiments, hot connection flow candidates may be identified based on the top N-tile of connection flows based on the maximum number of flows the CS may be expected to handle. For example, in at least one of the various embodiments, if a CS may be expected to handle 2000 connection flows, the top 25% of connection flows based on bit-rate (for a count of 500 flows) may identified as hot connection flows. Next, in at least one of the various embodiments, control may be returned to a calling process.

It will be understood that figures, and combinations of actions in the flowchart-like illustrations, can be implemented by computer program instructions. These program instructions may be provided to a processor to produce a machine, such that the instructions executing on the processor create a means for implementing the actions specified in the flowchart blocks. The computer program instructions may be executed by a processor to cause a series of operational actions to be performed by the processor to produce a computer implemented process for implementing the actions specified in the flowchart block or blocks. These program instructions may be stored on some type of machine readable storage media, such as processor readable non-transitive storage media, or the like.

What is claimed is:

1. A method for managing communication over a network with a traffic management device (TMD) that includes a plurality of components and is operative to perform actions, comprising:

employing at least one data flow segment (DFS) component to provide packet level flow handling for a portion of a plurality of connection flows:

employing at least one control segment (CS) component to perform actions, including:

managing the plurality of connection flows and handling a remainder portion of the plurality of connection flows:

generating at least one connection flow metric based on at least one received network packet for at least one of the plurality of managed connection flows;

employing the at least one connection flow metric to determine each hot connection flow in the plurality of managed connection flows;

- determining each hot connection flow to be handled by the DFS component wherein identifying each hot connection flow is based at least on a predicted connection flow capacity of the CS component, and wherein a percentile of connection flows are identified as hot connection flows; and
- employing the DFS component to handle each determined hot connection flow.
- 2. The method of claim 1, wherein employing the at least one connection flow metric further comprises determining 10 each hot connection flow handled by the DFS component if the plurality of managed connection flows exceeds a capacity of the DFS component.
- 3. The method of claim 1, wherein employing the at least one connection flow metric further comprises sorting the 15 plurality of managed connection flows based on at least a total amount of data exchanged over a time interval.
- **4.** The method of claim **1**, wherein employing the at least one connection flow metric further comprises:
  - determining a median bit-rate of data communicated for at 20 least one connection flow being handled by the CS component; and
  - employing the median bit-rate of data communicated for at least one connection flow and at least a bit-rate capacity of the CS component to estimate the maximum number 25 of connection flows the CS component can handle.
- 5. The method of claim 1, wherein determining each hot connection flow to be handled by the DFS component further comprises, identifying each hot connection flow to be handled by the DFS component based at least on a total 30 amount of data communicated over a time interval, wherein N number of connection flows having a top total amount of data communicated over the time interval are identified as hot connection flows.
- 6. The method of claim 1, wherein generating the at least 35 one connection flow metric further comprises, examining contents of the at least one received network packet to identify at least one of a data pattern, or a meta data which indicates that the at least one of the plurality of connection flows is a hot connection flow.
- 7. The method of claim 1, wherein the CS component performs further actions, comprising:
  - dividing each connection flow into an upload portion and a download portion;
  - generating separate connection flow metrics for each 45 upload portion and each download portion of each of the plurality of managed connection flows:
  - employing each connection flow metric to determine each hot download portion and each hot upload portion of the plurality of managed connection flows;
  - determining each hot upload portion and each download portion of the plurality of managed connection flows to be handled by the DFS component; and
  - employing the DFS component to handle each determined hot upload portion and each download portion of the 55 plurality of connection flows.
- **8**. A traffic management device (TMD) that includes a plurality of components for managing communication over a network and is operative to perform actions, comprising:
  - a transceiver that is operative to communicate data over the 60 network:
  - a memory that is operative to store instructions; and
  - a processor that is operative to execute instructions that enable actions, including:
  - employing at least one data flow segment (DFS) component to provide packet level flow handling for a portion of a plurality of connection flows; and

24

- employing at least one control segment (CS) component to perform actions, comprising:
  - managing the plurality of connection flows and handling a remainder portion of the plurality of connection flows;
  - generating at least one connection flow metric based on at least one received network packet for at least one of the plurality of managed connection flows;
  - employing the at least one connection flow metric to determine each hot connection flow in the plurality of managed connection flows;
  - determining each hot connection flow to be handled by the DFS component, wherein identifying each hot connection flow is based at least on a predicted connection flow capacity of the CS component, and wherein a percentile of connection flows are identified as hot connection flows; and
  - employing the DFS component to handle each determined hot connection flow.
- **9**. The TMD of claim **8**, wherein employing the at least one connection flow metric further comprises determining each hot connection flow handled by the DFS component if the plurality of managed connection flows exceeds a capacity of the DFS component.
- 10. The TMD of claim 8, wherein employing the at least one connection flow metric further comprises sorting the plurality of managed connection flows based on at least a total amount of data exchanged over a time interval.
- 11. The TMD of claim 8, wherein employing the at least one connection flow metric further comprises:
  - determining a median bit-rate of data communicated for at least one connection flow being handled by the CS component; and
  - employing the median bit-rate of data communicated for at least one connection flow and at least a bit-rate capacity of the CS component to estimate the maximum number of connection flows the CS component can handle.
- 12. The TMD of claim 8, wherein determining each hot connection flow to be handled by the DFS component further comprises, identifying each hot connection flow to be handled by the DFS component based at least on a total amount of data communicated over a time interval, wherein N number of connection flows having a top total amount of data communicated over the time interval are identified as hot connection flows.
  - 13. The TMD of claim 8, wherein generating the at least one connection flow metric further comprises, examining contents of the at least one received network packet to identify at least one of a data pattern, or a meta data which indicates that the at least one of the plurality of connection flows is a hot connection flow.
  - **14.** The TMD of claim **8**, wherein the CS component performs further actions, comprising:
    - dividing each connection flow into an upload portion and a download portion;
    - generating separate connection flow metrics for each upload portion and each download portion of each of the plurality of managed connection flows;
    - employing each connection flow metric to determine each hot download portion and each hot upload portion of the plurality of managed connection flows;
    - determining each hot upload portion and each download portion of the plurality of managed connection flows to be handled by the DFS component; and
    - employing the DFS component to handle each determined hot upload portion and each download portion of the plurality of connection flows.

- 15. A processor readable non-transitory storage media that is operative to store processor executable instructions for managing communication over a network with a traffic management device (TMD) having a plurality of components, wherein execution of the instructions by a processor enables 5 the TMD to perform actions, comprising:
  - employing at least one data flow segment (DFS) component to provide packet level flow handling for a portion of a plurality of connection flows;
  - employing at least one control segment (CS) component to 10 perform actions, including:
    - managing the plurality of connection flows and handling a remainder portion of the plurality of connection flows:
    - generating at least one connection flow metric based on 15 at least one received network packet for at least one of the plurality of managed connection flows;
    - employing the at least one connection flow metric to determine each hot connection flow in the plurality of managed connection flows;
    - determining each hot connection flow to be handled by the DFS component wherein identifying each hot connection flow is based at least on a predicted connection flow capacity of the CS component, and wherein a percentile of connection flows are identified as hot connection flows; and
    - employing the DFS component to handle each determined hot connection flow.
- **16**. The media of claim **15**, wherein employing the at least one connection flow metric further comprises determining 30 each hot connection flow handled by the DFS component if the plurality of managed connection flows exceeds a capacity of the DFS component.
- 17. The media of claim 15, wherein employing the at least one connection flow metric further comprises sorting the 35 plurality of managed connection flows based on at least a total amount of data exchanged over a time interval.
- 18. The media of claim 15, wherein employing the at least one connection flow metric further comprises:

26

- determining a median bit-rate of data communicated for at least one connection flow being handled by the CS component; and
- employing the median bit-rate of data communicated for at least one connection flow and at least a bit-rate capacity of the CS component to estimate the maximum number of connection flows the CS component can handle.
- 19. The media of claim 15, wherein determining each hot connection flow to be handled by the DFS component further comprises, identifying each hot connection flow to be handled by the DFS component based at least on a total amount of data communicated over a time interval, wherein N number of connection flows having a top total amount of data communicated over the time interval are identified as hot connection flows.
- 20. The media of claim 15, wherein generating the at least one connection flow metric further comprises, examining contents of the at least one received network packet to identify at least one of a data pattern, or a meta data which indicates that the at least one of the plurality of connection flows is a hot connection flow.
- 21. The media of claim 15, wherein the CS component performs further actions, comprising:
  - dividing each connection flow into an upload portion and a download portion;
  - generating separate connection flow metrics for each upload portion and each download portion of each of the plurality of managed connection flows;
  - employing each connection flow metric to determine each hot download portion and each hot upload portion of the plurality of managed connection flows;
  - determining each hot upload portion and each download portion of the plurality of managed connection flows to be handled by the DFS component; and
  - employing the DFS component to handle each determined hot upload portion and each download portion of the plurality of connection flows.

\* \* \* \* \*